

Balls into Bins with Strict Capacities and Weighted Edges

Ankur Sahai

No Institute Given

Abstract. We explore a novel theoretical model for studying the performance of distributed storage management systems where the data-centers have limited capacities (as compared to storage space requested by the users). Prior schemes such as Balls-into-bins (used for load balancing) neither consider bin (consumer) capacities (multiple balls into a bin) nor the future performance of the system after, balls (producer requests) are allocated to bins and restrict number of balls as a function of the number of bins. Our problem consists of finding an optimal assignment of the online producer requests to consumers (via weighted edges) in a complete bipartite graph while ensuring that the total size of request assigned on a consumer is limited by its capacity. The metric used to measure the performance in this model is the (minimization of) weighted sum of the requests assigned on the edges (loads) and their corresponding weights. We first explore the optimal offline algorithms followed by the analysis of different online techniques (by comparing their performance against the optimal offline solution). LP and Primal-Dual algorithms are used for calculating the optimal offline solution in $O(r \cdot n)$ time (where r and n are the number of requests and consumers respectively) while randomized algorithms are used for the online case.

We propose randomized online algorithms in which the consumers are selected based on edge probabilities (that can change with consumer failures; due to capacity exhaustion) and evaluate the performance of these randomized schemes using probabilistic analysis. The performance of the online algorithms is measured using competitive analysis assuming an oblivious adversary who knows the randomized algorithm but not the results produced. For the simplified model with equal consumer capacities an average-case competitive ratio (which compares the average cost of the output produced by the online algorithm and the minimum cost of the optimal offline solution) of $(\frac{d}{\min_{i,j} d_{i,j}})$ (where d is the edge weight / distance) is achieved using an algorithm that has equal probability for selecting any of the available edges with a running time of $O(r)$. In the extending the model to arbitrary consumer capacities we show an average case competitive ratio of $(\frac{d \cdot c}{\bar{c} \cdot \min_{i,j} d_{i,j}})$. This theoretical model gives insights to a (storage) cloud system designer about, how the different attributes (producer requests, edge weights and consumer capacities) effect the overall (read / write) performance of a distributed storage management system over a period of time.

Keywords: Online algorithms, Primal-dual, Load balancing, Competitive ratio.

1 Motivation

A (storage cloud) is a collection of VMs (represented by consumers in our model in section 2) and data-centers (producers) that interact with each other through network links with different bandwidths / data-transfer rates (inversely proportional to the edge weights). VMs generate requests for virtual disk space (requests) at the time of creation. The data-centers have limited storage space (capacity). One of the core problems is to optimally select a data-center for allocating the disk space for a VM. If we assume that all data-centers are available for the requests generated by each VM then this configuration forms a complete bipartite graph. As the VM users are added gradually to the cloud and their disk-space requirements need to be satisfied instantly to meet the SLAs this forms the online part of the problem. We assume that once a VM is allocated storage space on a particular data-center it cannot be moved to a different one, thus VMs read / write data using the same network links (fixed at the time of VM provisioning) throughout their lifetime.

The overall I/O performance of the cloud storage system can be optimized when the quality of network links used for the majority of read / write operations is maximized which, can be achieved by using the higher quality (inversely proportional to edge weights) links for allocating as much producer requests as possible. This is equivalent to minimizing the the distance weighted sum of request allocated in the system. This theoretical model can be used to measure the performance of distributed storage provisioning schemes such as VMware's virtualization framework - VSphere [8,9].

2 Problem Definition

In the offline model \mathcal{M}_{OFF} the requests sizes $s(t)$, capacities c_j and edge weights $d_{i,j}$ are arbitrary.

- A complete (undirected) bipartite graph $G = (P, C, E)$.
- Set $P = \{P_i \mid i \in [m]\}$ of producers, $m = |P|$. Set $C = \{C_j \mid j \in [n]\}$ of consumers, $n = |C|$, with arbitrary capacity c . The capacity is a strict bound for the consumer's load and cannot be exceeded.
- Set of edges $E = \{e_{i,j} \mid i \in [m], j \in [n]\}$ connect each producer to each consumer and with $d_{i,j}$ which denote the distance between producer P_i and C_j .
- r is the total number of requests. The requests are produced by the producers and allocated to the consumers. A request can be split and allocated to different producers. $s(t)$ is the size of the request produced at time t .
- Producers are selected at random in each round.
- $d(t)$ is the distance between the producer and consumer that are chosen in time step t . The total cost of the random allocation is $G = \sum_{t=1}^r s(t) \cdot d(t)$.
- We assume that every request can commit to at least one consumer (regardless of the previous allocation), i. e. at least one consumer must have sufficient space left. The number and size of requests must be limited accordingly.
- ℓ_j is the load on consumer C_j . All loads ℓ_1, \dots, ℓ_n are set to 0 in the beginning. Whenever a consumer receives a request, its load is increased by the size of the request.

- $\ell_{i,j}$ is the total request from producer P_i allocated to consumer C_j known as the edge load. All loads $\ell_{1,1}, \dots, \ell_{m,n}$ are set to 0 in the beginning. Whenever a consumer receives a request from a producer the corresponding edge load is increased by the size of the request.

The objective of this problem is to find an average-case [2] α -competitive online algorithm (assuming startup cost $\beta = 0$), that minimizes the expected value of weighted sum of edge loads while satisfying the producer requests:

$$\sum_{i \in [m], j \in [n]} E_I \left[d_{i,j}(t) \cdot l_{i,j}(t) \right] \leq \alpha \cdot OPT_I(t), \forall(t, I) \quad (1)$$

where, α is a constant and $OPT_I(t)$ is the output of the optimal offline algorithm for the input received in the time interval $(0, t]$ for an instance I .

Equation (2) guarantees that the total load on the edges $(\sum_{j \in [n]} l_{i,j}(t))$ incident on each producer $(i \in [m])$ is equal to the total size of the requests generated by the producer $(\sum_t R_i(t))$. This will be referred to as *producer request constraint*.

$$\sum_{j \in [n]} l_{i,j}(t) = \sum_t R_i(t), \forall(i \in [m], t) \quad (2)$$

Equation (3) ensures that the total load on the edges $(\sum_{i \in [m]} l_{i,j}(t))$ incident on a consumer $(j \in [n])$ does not exceed the consumer capacity (c_j) . This will be referred to as *consumer capacity constraint*.

$$\sum_{i \in [m]} l_{i,j}(t) \leq c_j, \forall(j \in [n], t) \quad (3)$$

3 Related Work

Balls-into-Bins [1,4,6] model is used for studying load balancing in a similar resource allocation configuration where, the objective is to place m balls into n bins while guaranteeing bounds on the maximum, minimum or the average load across all the bins. The main advantage of the model defined in section 2 is that it also takes into accounts the capacities of the bins (which are analogous to the consumers in the problem defined in section 2). Further the model described in section 2 compares the performance of a randomized algorithm with the optimal offline algorithm.

One of the well studied online algorithm for assigning resources to users is the k -server problem [7] where, the servers handle request once for each client. Dynamic assignment [3] has a similar configuration involving bipartite graphs.

4 Offline Algorithms

The optimal offline algorithm has to exhaustively look at the available edges for allocating a request. This paper uses Linear Programming and Primal-Dual algorithms for solving the offline version (section 2).

4.1 Linear Programming

Linear Programming [5] is a method used to solve large-scale optimization problems with a set of constraints and an objective function (minimization or maximization) both being linear.

The LP formulation for this problem is as follows,

Objective function:

Minimize:

$$\sum_{i \in [m], j \in [n]} d_{i,j}(t) \cdot l_{i,j}(t), \quad d_{i,j}(t) \geq 0, \quad l_{i,j}(t) \geq 0 \quad (4)$$

Constraints:

$$\sum_{j \in [n]} l_{i,j}(t) \geq \sum_t R_i(t), \quad R_i(t) > 0, \quad \forall(i \in [m], t) \quad (5)$$

$$\sum_{i \in [m]} l_{i,j}(t) \leq c_j \implies - \sum_{i \in [m]} l_{i,j}(t) \geq -c_j, \quad \forall(j \in [n] t) \quad (6)$$

This LP is used for calculating the optimal solution $\text{OPT}(t)$ for the input received in the time interval $(0, t], t$. As the requests $R_i, i \in [m]$ are non-negative the load assignments $l_{i,j}$ in the objective function (4) are also non-negative. Equation (5) represents the *producer request* constraint corresponding to (2) whereas (6) corresponds to the *consumer capacity* constraint corresponding to (3). In addition to this, new constraints corresponding to the existing load assignment on edges have to be added at each time instance t for *assignment without reallocation* (defined in section 2).

LP formulation in section 4.1 produces a feasible (without violating the consumer capacity constraints (3)) assignment of loads $l_{i,j}$ on edges $e_{i,j}$ corresponding to the requests $R_i(t), \forall(i \in [m], t)$. Equation (5) guarantees that the total request generated by producers $i \in [m]$ is satisfied and (6) ensures that the capacities of consumers $j \in [n]$ are not exceeded. By definition of problem in section 2, this is a valid assignment of loads on edges. LP formulation in section 4.1 produces the optimal assignment of loads $l_{i,j}$ on edges $e_{i,j}$ corresponding to the requests $R_i(t), \forall(i \in [m], t)$. The solution produced by LP is optimal as fractional loads are allowed and the LP considers all possible solutions minimize the objective function (4).

4.2 Primal-Dual

Primal-Dual algorithms are used for a certain class of optimization problems where there are a finite number of feasible solutions available at each step. The dual is often useful for providing intuitions about the nature of the solution that are implicit in the primal.

Consider the dual of the LP formulation in section 4.1. Let y_i be the dual variables corresponding to producers $i \in [m]$ (5) and z_j be the dual variables corresponding to the consumers $j \in [n]$ (6) then the corresponding dual is,

Objective function:

Maximize:

$$\sum_{i \in [m]} y_i \cdot R_i - \sum_{j \in [n]} z_j \cdot c_j, y_i \geq 0, z_j \geq 0 \quad (7)$$

$R_i = \sum_{t=1}^T R_i(t)$ is the total size of the requests produced by P_i

Constraints:

$$y_i - z_j \leq d_{i,j}, \forall (i \in [m], j \in [n]) \quad (8)$$

Equation (8) suggests that the potential difference between producers and consumers can be at most equal to $d_{i,j}$. This will be referred to as dual *potential limit* constraint.

According to complementary slackness conditions,

$$l_{i,j} > 0 \iff y_i - z_j = d_{i,j} \quad (9)$$

Let T be the set of tight constraints (which represent edges selected by Algorithm 1) ($T = \{(i, j) \mid y_i - z_j = d_{i,j}, \forall (i \in [m], j \in [n])\}$) and S , the set of slack constraints ($S = \{(i, j) \mid y_i - z_j < d_{i,j}, \forall (i \in [m], j \in [n])\}$). Then consider an *unit benefit* function which measures the increase in dual objective (7) corresponding to a unit amount of request ($\Delta y_i \cdot R_i(t)$) produced by producer $i \in [m]$ at time t where, Δz_j are the corresponding increases in the consumer-dual variables; required for keeping the dual *potential limit* constraints (8) tight.

$$B_i(t) = \Delta y_i \cdot R_i(t) - \sum_{j \in T} \Delta z_j \cdot c_j \quad (10)$$

The *unit benefit* function in 10 is used as a criteria for selecting the request which produces the maximum increase in the dual objective function by the Primal-dual algorithm.

Algorithm 1 Primal-Dual algorithm for *Offline Assignment*

```

1:  $T \leftarrow \emptyset$ 
2:  $S \leftarrow \{(i, j) \mid \forall (i \in [m], j \in [n])\}$ 
3:  $y_i \leftarrow 0, \forall i \in [m]$ 
4:  $z_j \leftarrow 0, \forall j \in [n]$ 
5:  $l_{i,j} \leftarrow 0, \forall (i \in [m], j \in [n])$ 
6: while  $\exists (i, t) : B_i(t) \geq 0$  do
7:    $y_i : \max_i (B_i(t))$ 
8:    $z_j : \min_j (d_{i,j} - (y_i - z_j) \mid i \in [m], j \in [n])$ 
9:    $\Delta_1 = d_{i,j} - (y_i - z_j)$ 
10:   $y_i \leftarrow y_i + \Delta_1$ 
11:   $z_j \leftarrow z_j + \Delta_1, \forall j \in T$ 
12:   $T = T \cup (i, j)$ 
13:   $S = S \setminus (i, j)$ 
14:   $\Delta_2 = c_j - \sum_{(i,j) \in T} l_{i,j}$ 
15:   $l_{i,j} \leftarrow l_{i,j} + \Delta_2$ 
16: end while

```

The primal-dual Algorithm 1 chooses the (producer dual variable y_i corresponding to) request $R_i(t)$ with the highest benefit function $B_i(t)$ (step 7 in Algorithm 1) to maximize the increase in value of dual objective function and then chooses the dual *potential limit* constraint (operation 12) (and the corresponding variable z_j) that is closest to becoming tight at step 8 in Algorithm 1 (corresponding to the least cost edge in the primal) and increases the value of y_i by the amount (Δ_1 , in steps 9 and 10 in Algorithm 1) that is needed to make this constraint tight (which corresponds to selecting an edge for allocating a request in the primal using complementary slackness condition 9). If this constraint is tight (T updated in step 12) then corresponding z_j are also increased by Δ_1 (step 11 in Algorithm 1) to maintain tightness.

Increasing $d_{i,j}$ by amount Δ_2 is only symbolic / superficial. The practical output of the algorithm can be traced by looking up the set of dual constraints (corresponding to edges in primal) that are selected corresponding to each request. Intuitively, optimal offline primal-dual Algorithm 1 assigns the lower sized requests to the higher cost edges and uses the lower cost edges for the higher sized requests.

Theorem 1 (Optimality of Algorithm 1). *The Primal-Dual Algorithm 1 reaches the optimal solution for the assignment without reallocation problem in section 2 when it is not possible to increase the cost of the dual objective function (7).*

Proof. initializing $y_i \leftarrow 0, \forall i \in [m]$ and $z_j \leftarrow 0, \forall j \in [n]$ (steps 1 and 2 in Algorithm 1) and from the way we increase the dual variables (step 9 in Algorithm 1) the dual *potential limit constraint* (8) is always satisfied. Setting $l_{i,j} \leftarrow 0, \forall (i \in [m], j \in [n])$ (step 5 in Algorithm 1) makes sure that the primal is feasible at the start. The primal *consumer capacity* constraints in (6) are not violated by the way we increase $l_{i,j}$ from steps 14 and 15 in Algorithm 1.

When the cost of the dual objective function (7) cannot be increased further,

$$\sum_{i \in T} \Delta y_i \cdot R_i - \sum_{j \in T} \Delta z_j \cdot c_j = 0 \quad (11)$$

$$\implies \sum_{i \in T} \Delta y_i \cdot R_i = \sum_{j \in T} \Delta z_j \cdot c_j \quad (12)$$

At this point, if there was a pending request $R_i(t)$ with a positive *unit benefit* function that could choose one of the slack (S) dual constraints (without violating any of the *potential limit* dual constraints(8)) then the corresponding producer dual variable y_i can be increased until the dual constraint became tight. By contradiction such a request does not exists by definition (step 6) of Algorithm 1 otherwise, the algorithm would have continued. So the system is in equilibrium and changing the value of any of the dual variables will violate (one or more) dual *potential limit* constraints.

However, for the current solution $((y_i, z_j) \in T)$, the dual variables need to be changed by an equal amount ($\Delta y_i = \Delta z_j = \Delta, \forall (i \in [m], j \in [n])$) for maintaining the potential equilibrium [keeping the dual constraints in T tight for meeting complementary slackness (9) condition].

$$\Delta y_i = \Delta z_j = \Delta \quad (13)$$

When it is not possible to increase the value of dual objective function then using (12) and (13),

$$\sum_{i \in T} R_i = \sum_{j \in T} c_j \quad (14)$$

From steps 14 and 15 in Algorithm 1, we know that,

$$\sum_{j \in T} c_j = \sum_{(i,j) \in T} l_{i,j}(t) \quad (15)$$

Using (14) and (15) we infer that the requests have been met,

$$\sum_{i \in T} R_i = \sum_{j \in T} l_{i,j}(t) \quad (16)$$

This means that the primal *producer request* constraints in (5) are feasible and the primal *consumer capacity constraints* in (6) and dual *potential limit* constraints are always feasible. Complementary slackness (9) is satisfied as we only increase $d_{i,j}$ when the dual constraint is tight. This means that the primal is optimal and the dual is optimal. Thus, this is the optimal solution for the problem in section 2.

Theorem 2 (Time complexity of Algorithm 1). *The Primal-Dual Algorithm 1 takes $O(r \cdot n)$, time to complete where, d and n the number of requests and consumers respectively.*

Proof. At each step of the Algorithm 1 $O(r)$ (where r is the number of requests) operations are needed to select the (producer dual variable y_i corresponding to the) request with maximum *unit benefit* function, another $O(n)$ operations to find the corresponding consumer dual variable z_j . Hence, it takes $O(d + n)$ operations for the first (while loop from steps 6 to 16 in Algorithm 1) and $O(d - 1 + n)$ operations for the second request by keeping track of the requests that have been covered by setting a flag in a hash table. This gives a time complexity of $O(r \cdot n)$ using amortized analysis.

5 Online Algorithms

In the beginning we look at the following simplified model \mathcal{M} ,

- Let $R(t)$ be the request in round t and let $C(t)$ denote the set of consumers in round t that have enough space left to store request $R(t)$.
- All requests $R(t)$ have the same size s .
- All consumers C_i have equal capacities c .
- An oblivious online adversary who can only manipulate the size of the online requests.

We use randomized algorithms for the online version as for a deterministic algorithm the oblivious online adversary who knows the edge selected in each round can manipulate the online request sizes to distort the value of the objective function (1).

The simplified algorithm \mathcal{A} is:

Algorithm 2 Online Algorithm \mathcal{A}

```
1: Given:  $P; C; c; d_{i,j}; s$ 
2: Initialize:  $\ell_j = 0 \forall j \in [n]; G = 0; n = |C|$ 
3: for  $t = 1$  to  $r$  do
4:   Choose  $P_i \in [m]$  independently and uniformly at random
5:    $C(t) = \{c_j \mid j \in [n] \wedge \ell_j + s \leq c\}$ 
6:   Choose  $c_j \in [n]$  independently and uniformly at random from  $C(t)$ .
7:    $\ell_j = \ell_j + s$ 
8:    $G = G + s \cdot d_{i,j}$ 
9: end for
```

Theorem 3. Assuming model \mathcal{M} and running algorithm \mathcal{A} , the expected total cost is $E[G] = \sum_{t=1}^r s \cdot \bar{d}$.

Proof. Using the law of total probability over the different combination of consumer failures $|C_j| = n - k$ when there are $n - k$ consumers available,

$$p(e_{i=P_i,j} \mid (|C(t)| = n - k)) = \sum_{j=1}^{\binom{n}{k}} p(e_{i=P_i,j} \mid |C_j| = n - k) \cdot p(|C_j| = n - k) \quad (17)$$

$$= \frac{1}{\binom{n}{k}} \left(\frac{\binom{n}{k} \cdot \sum_{j=1}^n p(e_{i,j}) - \frac{\binom{n}{k} \cdot k}{n} \cdot \sum_{j=1}^n p(e_{i,j})}{n - 1} \right) \quad (18)$$

$$= \frac{\sum_{j=1}^n d_{i,j}}{n} \quad (19)$$

Using (19) and the law of total expectation over the number of consumers available,

$$E[d_{i=P_i,j} \mid P_i] = \sum_{k=1}^n E[d_{i=P_i,j} \mid (|C(t)| = n - k)] \cdot p(|C(t)| = n - k) \quad (20)$$

$$= \frac{\sum_{j=1}^n d_{i,j}}{n} \cdot \sum_{k=1}^n p(|C(t)| = n - k) \quad (21)$$

$$= \frac{\sum_{j=1}^n d_{i,j}}{n} \quad (22)$$

We calculate the expected distance of an edge selected in any round over the scenarios corresponding to different number k of consumers available using (21). As the probability to select an edge does not change in the different scenarios we deduce in (18), (19) and (20) that, there is uniform probability of selecting any (consumer) edge given that producer P_i is selected in a round. Intuitively this follows from symmetry as each consumer is equally likely to be picked in any round assuming equal capacities.

Using the law of total expectation over the producer selected in a round,

$$E[d_{i,j}] = \sum_{i=1}^m E[d_{i=P_i,j} \mid i = P_i] \cdot p(i = P_i) \quad (23)$$

$$= \sum_{i=1}^m \left[\frac{1}{n} \cdot \sum_{j=1}^n d_{i,j} \right] \cdot \frac{1}{m} \quad (24)$$

$$= \frac{1}{m \cdot n} \sum_{i=1}^m \sum_{j=1}^n d_{i,j} = \bar{d} \quad (25)$$

So the expected cost of Algorithm \mathcal{A} (2) in any round is $E[G] = s \cdot E[d_{i,j}] = s \cdot \bar{d}$. The best possible cost of the optimal offline algorithm OPT in any round is when it chooses the cheapest edge $\min_{i,j} d_{i,j}$ in each round. Thus, the average-case ([2]) competitive ratio is $\left(\frac{\bar{d}}{\min_{i,j} d_{i,j}} \right)$. The worst-case competitive ratio is $\left(\frac{d_{\max}}{\min_{i,j} d_{i,j}} \right)$.

Theorem 4. *Extending model \mathcal{M} to $\mathcal{M}2$ with arbitrary producer request sizes, equal consumer capacities, arbitrary edge distances and allowing a single requests to be allocated across multiple consumers and running algorithm $\mathcal{A}2$, the expected total cost is $E[G] = \sum_{t=1}^r s(t) \cdot \bar{d}$.*

Proof. The simplified algorithm $\mathcal{A}2$ is:

Algorithm 3 Online Algorithm $\mathcal{A}2$

- 1: Given: $P; C; c; d_{i,j}; s$
 - 2: Initialize: $\ell_j = 0 \forall j \in [n]; G = 0; n = |C|$
 - 3: **for** $t = 1$ to r **do**
 - 4: Choose $P_i \in [m]$ independently and uniformly at random
 - 5: Choose an arbitrary request size $s(t) = \left(0, \sum_{j \in C} (c_j - \ell_j) \right]$
 - 6: **for** $s = 1$ to $s(t)$ **do**
 - 7: $C = \{c_j \mid j \in [n] \wedge \ell_j + s(t) \leq c\}$
 - 8: Choose $c_j \in [n]$ independently and uniformly at random
 - 9: $\ell_j = \ell_j + 1$
 - 10: $G = G + d_{i,j}$
 - 11: **end for**
 - 12: **end for**
-

By splitting the requests into unit sized blocks of size f_D and using the fact that there is equal probability for the blocks to be allocated to any consumer, it follows from symmetry that a block of unit requests is equally likely to belong to any producer P_i . Using (20), (22) and (26) from **Theorem 3** we get the expected cost of the edge selected in any round as \bar{d} . So the expected cost is $E[G] = \sum_{t=1}^r s(t) \cdot \bar{d}$ and the average case competitive ratio is $\left(\frac{\bar{d}}{\min_{i,j} d_{i,j}} \right)$

Theorem 5. Extending model M2 to M3 with arbitrary producer demands, arbitrary consumer capacities and arbitrary edge distances and running algorithm A3, the expected total cost is $E[G] = \sum_{t=1}^r s(t) \cdot \frac{d_{\text{avg}}}{c}, \forall t$.

Proof. The simplified algorithm A3 is,

Algorithm 4 Online Algorithm A3

```

1: Given:  $P; C; c; d_{i,j}; s$ 
2: Initialize:  $\ell_j = 0 \forall j \in [n]; G = 0; n = |C|$ 
3: for  $t = 1$  to  $r$  do
4:   Choose  $P_i \in [m]$  independently and uniformly at random
5:   Choose an arbitrary request size  $s(t) = \left(0, \sum_{j \in C} (c_j - \ell_j)\right]$ 
6:   for  $s = 1$  to  $s(t)$  do
7:      $C = \{c_j \mid j \in [n] \wedge \ell_j + s(t) \leq c_j\}$ 
8:     Choose  $c_j \in n$  with probability  $\left(\frac{c_j}{\sum_{j=1}^n c_j}\right)$ 
9:      $\ell_j = \ell_j + 1$ 
10:     $G = G + d_{i,j}$ 
11:   end for
12: end for

```

Let $X_{i,j}$ be a indicator random variable that indicates that producer P_i selects consumer C_j . As requests are split into unit sized blocks, the expected load on consumer C_j assigned by producer P_i is equal to value of $X_{i,j}$. Then $E[X_{i,j}] = E[C_j \mid P_i] = \frac{c_j}{\sum_{j \in n} c_j}$. Let Y_j be the indicator random variable for the load placed on consumer C_j in a specific round then, $E[Y_j] = \sum_{P_i=1}^m E[X_{P_i,j}] \cdot p(P_i) = \frac{c_j}{\sum_{j \in n} c_j}$ as each producer is equally likely to be selected in a round. Let $Y_j(t)$ denote the load placed on consumer C_j after r requests have been completed then $E[Y_j(t)] = \sum_{t=1}^r s(t) \cdot E[Y_j] = \frac{c_j}{\sum_{j \in [n]} c_j} \cdot \sum_{t=1}^r s(t)$. Using unit sized requests the consumer load always remains proportional to its capacity.

As the producers are chosen uniformly at random the probability that a consumer C_j was picked by a producer P_i in any round is $p[P_i \mid C_j] = \frac{1}{m}$. Size of the request chosen $(0, \sum_{j \in C} (c_j - \ell_j)]$ at each iteration can be atmost equal to the remaining capacity available as it can be split amongst the consumers.

Expected distance of the edge when consumer C_j is picked is,

$$E[d_{ij} \mid C_j] = \sum_{i=1}^m d_{i,j} \cdot p(P_i \mid C_j) \quad (26)$$

$$= \frac{1}{m} \cdot \sum_{i=1}^m d_{i,j} \quad (27)$$

The expected cost of edge picked in a any round assuming equal failure times,

$$E[d_{i,j}] = \sum_{j=1}^n E[d_{i,j} \mid c_j] \cdot p(c_j) \quad (28)$$

$$= \frac{1}{m} \cdot \sum_{j=1}^n \sum_{i=1}^m d_{i,j} \cdot \frac{c_j}{\sum_{j=1}^n c_j} \quad (29)$$

$$= \frac{n}{\sum_{j=1}^n c_j} \cdot \frac{\sum_{j=1}^n \sum_{i=1}^m d_{i,j} \cdot c_j}{m \cdot n} = \frac{\overline{d \cdot c}}{\overline{c}} \quad (30)$$

Expected cost is $E[G] = \sum_{t=1}^r s(t) \cdot \frac{\overline{d \cdot c}}{\overline{c}}$ and the average-case competitive ratio is $\left(\frac{\overline{d \cdot c}}{\overline{c} \cdot \min_{i,j} d_{i,j}} \right)$. Although the competitive ratio depends on the edge distances and consumer capacities, it is not possible to assume equal failure times using any other edge probability although this online algorithm $\mathcal{A}3$ is not optimal for the objective function in (1).




Fig. 1. Comparison of Greedy and Randomized Vs. Optimal Offline for special cases

The instances considered ranged from 1 to 100 producers and 1 to 100 consumers. The size of demands, edge distances and consumer capacities were picked randomly.

6 Conclusion

The average-case competitive ratio of $\left(\frac{\bar{d}}{\min_{i,j} d_{i,j}}\right)$ for the model with equal consumer capacities and arbitrary producer requests indicates that the performance depends on the quality of majority of links. For arbitrary consumer capacities with an average-case competitive ratio of $\left(\frac{\bar{d} \cdot c}{c \cdot \min_{i,j} d_{i,j}}\right)$ the performance is decided by the quality of the links connected to the consumers with higher capacities. The optimal offline primal-dual algorithm runs in $O(r \cdot n)$ time whereas the online algorithms take $O(r)$ time where r and n are the number of requests and consumers respectively.

References

1. Raab, C Steger, A, "Balls into bins" - A simple and tight analysis, Randomization and Approximation Techniques in Computer Science, LNCS, **1518**, 159-170, 1998
2. Fujiwara, H, Iwama, K, Bose, P, Morin, P Average-Case Competitive Analyses for Ski-Rental Problems Algorithms and Computation, LNCS, Springer Berlin / Heidelberg, 2002, **2518**, 157-189.
3. Mills-Tettey, G. A., Stentz, A. T., and Dias, C. B., The dynamic Hungarian algorithm for the assignment problem with changing costs, Tech. rep. CMU-RI-TR-07-27, Robotics Institute, 2007.
4. K. Schrder, Balls into Bins: A Paradigm for Job Allocation, Data Distribution Processes, and Routing, PhD Thesis, 2000.
5. Correa JR, Wagner MR, LP-based online scheduling: From single to parallel machines, Integer programming and combinatorial optimization, proceedings, lecture notes in computer science, **3509**, 196-209, 2005.
6. Azar Y, Broder AZ, Karlin AR, Upfal E, Balanced allocations, SIAM Journal of Computing, **29**,1, 180-200, Sep 1999.
7. Yair Bartal, Manor Mendel, Randomized k-server algorithms for growth-rate bounded graphs, Journal of Algorithms **55** (2005) 192202.
8. Dynamic Storage Provisioning, VMware, Nov 18, 2009.
9. Scalable Storage Performance, VMware, Jun 5, 2008.